NOνA Answers to NuSAG Questions

July 27, 2005

Question 1: What is the assumed systematic error on the background in your sensitivity projections? What measurements, yours and other experiments', are required to reach this assumed error? How are these measurements to be performed?

Answer to Question 1: We have assumed a 5% systematic error on the background in all of our physics projections. We have not yet done a careful study of our ultimate error on backgrounds, although such a study is high on our priority list. In the following, we outline the steps that an analysis would go through to determine the backgrounds.

NOνA is a two-detector experiment. The backgrounds in the Far Detector are scaled from the measurements in the Near Detector. This permits, to a great extent, a cancellation of errors due to uncertainty in cross sections and detection efficiency.

The major problem in extrapolating the backgrounds in the Near Detector to the Far Detector is that the different components of the background scale differently. However, the sources of these differences are well understood, and can be calculated with reasonable precision. First, the Near Detector spectra are broader than the Far Detector spectra because the Near Detector sees a line source while the Far Detector effectively sees a point source. Second, the beam ν_e backgrounds scale differently than the ν_μ neutral current (NC) backgrounds because the ν_e backgrounds come primarily from muon decay, which occurs on the average further downstream in the decay pipe. Third, the ν_μ charged current (CC) backgrounds scale differently because they tend to oscillate away at the Far Detector. And finally, there will be a rather small background contribution from ν_τ CC events, which needs to be included.

Given this, what is the process that we need to go through to determine the proper extrapolation of backgrounds? The first step would be to update the Monte Carlo beam simulations with the most up-to-date information, mainly from MIPP (see answer to Question 7) and from MINOS Near Detector measurements. The latter will give expected ν_μ rates in the NOνA Near Detector, since same pion gives different energy neutrinos into detectors on and off-axis, and the ν_μ CC cross sections at different energies will have been measured both by the MINERνA experiment and by the MINOS Near Detector. Subsequently one analyzes the NOνA Near Detector data to determine the rate of ν_e CC, ν_μ CC and NC events, taking into account misidentifications. These misidentifications are calculated on the basis of Monte Carlo simulations, which will incorporate latest information on neutrino interactions from the MINOS Near Detector, from the NOνA Prototype Near Detector, and from the MINERνA experiment. The Monte Carlo can then be more finely tuned using the data from the NOνA Near Detector in different positions, as discussed below. The final Monte Carlo is then used to extrapolate the Near Detector spectra to the expected Far Detector spectra without any oscillations. The calculation of backgrounds to $\nu_\mu \rightarrow \nu_e$ oscillations requires modifying the extrapolated spectra for the effects of ν_μ CC disappearance, as measured by MINOS and eventually NOνA (discussed in the answer to Question 3).

We will be aided in constraining the extrapolation factors by being able to move the Near Detector in the NuMI tunnel. For example, Fig. 1.1 shows that at Site 1.5, there will be very good agreement between Near and Far detectors for the ν_e CC spectrum, while Fig. 1.2 shows that the ν_μ CC (and NC) spectra will be best represented by a location somewhere between Sites 1.5 and 2. Figure 1.3 shows that this location is Site 1.75. The Near Detector spectrum is wider than the Far Detector spectrum, for the reason given above, but the peaks match perfectly. The good agreement between the ν_e spectra at Site 1.5 and the ν_μ spectra at Site 1.75 give us confidence that we will be able to perform the near to far extrapolation quite precisely. Note that there should be little model dependence in these extrapolations since the MIPP data (discussed in more detail in the answer to Question 7) will determine the neutrino fluxes at the Near and Far Detector sites to the accuracy of these measurements.

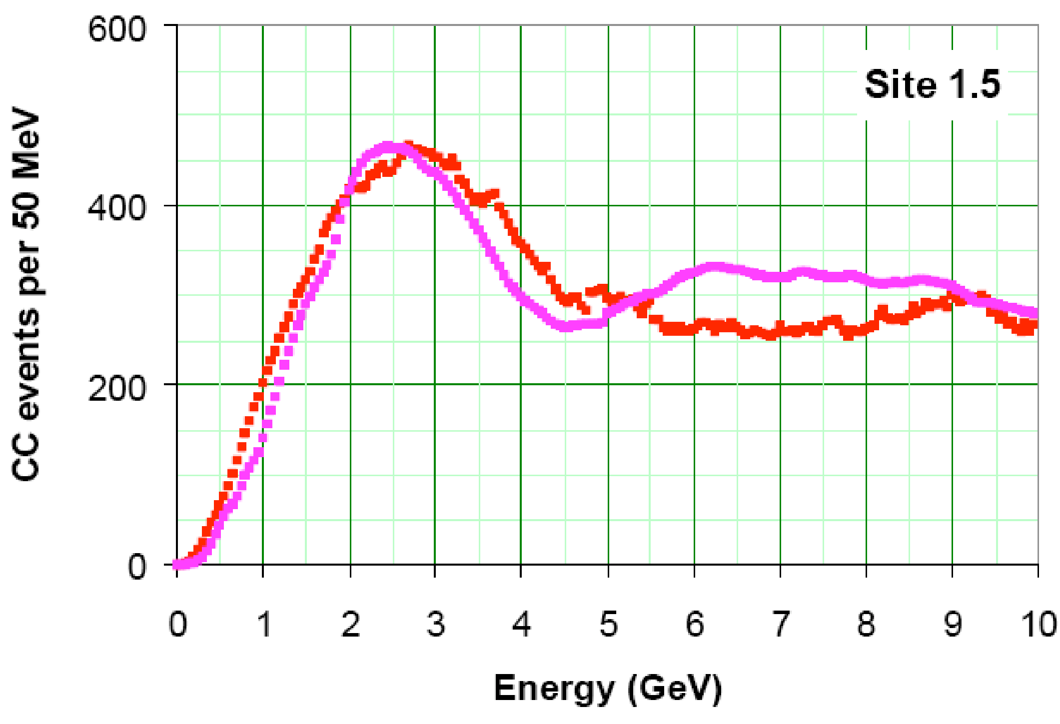

Fig. 1.1: The beam ν_e event rates for one year of data in the Near Detector located at Site 1.5 (red squares). The Far Detector beam ν_e distribution (pink line) is also shown assuming no oscillation, but has been normalized to have the same value at 2 GeV.

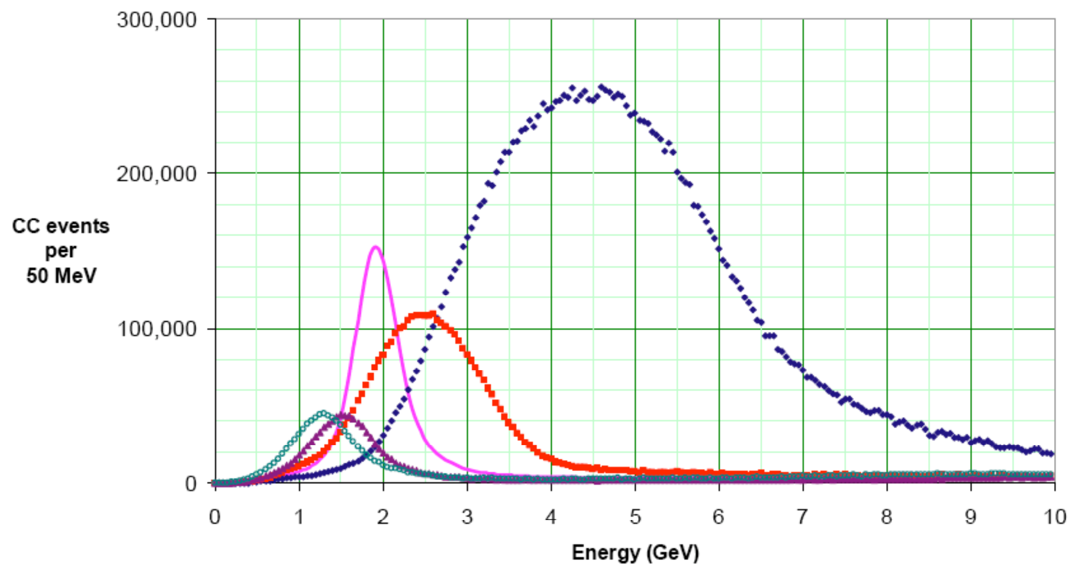

Fig. 1.2: ν_μ CC spectra for the various Near Sites [Site 1(blue diamonds), Site 1.5 (red squares), Site 2 (purple triangles), Site 3 (open green circles)] for one year of running at 6.5 10^{20} pot. The unoscillated Far Detector spectrum for one year of running (times an arbitrary scale factor of 800) is shown as the solid pink line.

In general, the method of modifying the Monte Carlo to agree with the Near Detector will not be unique, and the different backgrounds at the Far Detector that result from the different methods of modifying the Monte Carlo will result in systematic uncertainties. For example, one of the larger uncertainties might be the determination of the fraction of apparent Near Detector NC events that are actually CC events in which the muon is not detected. The systematic uncertainties only start becoming significant after a couple of years of running (see the answer to Question 2), so we will have adequate data and time to optimize the Monte Carlo.

A detailed study of the extrapolation process will allow us to evaluate how accurately we can determine these backgrounds, and how valuable measurements from the MIPP and MINERνA experiments will be in this determination. It will also allow us to determine the value of our plan to move the Near Detector to different locations along the NuMI tunnel.

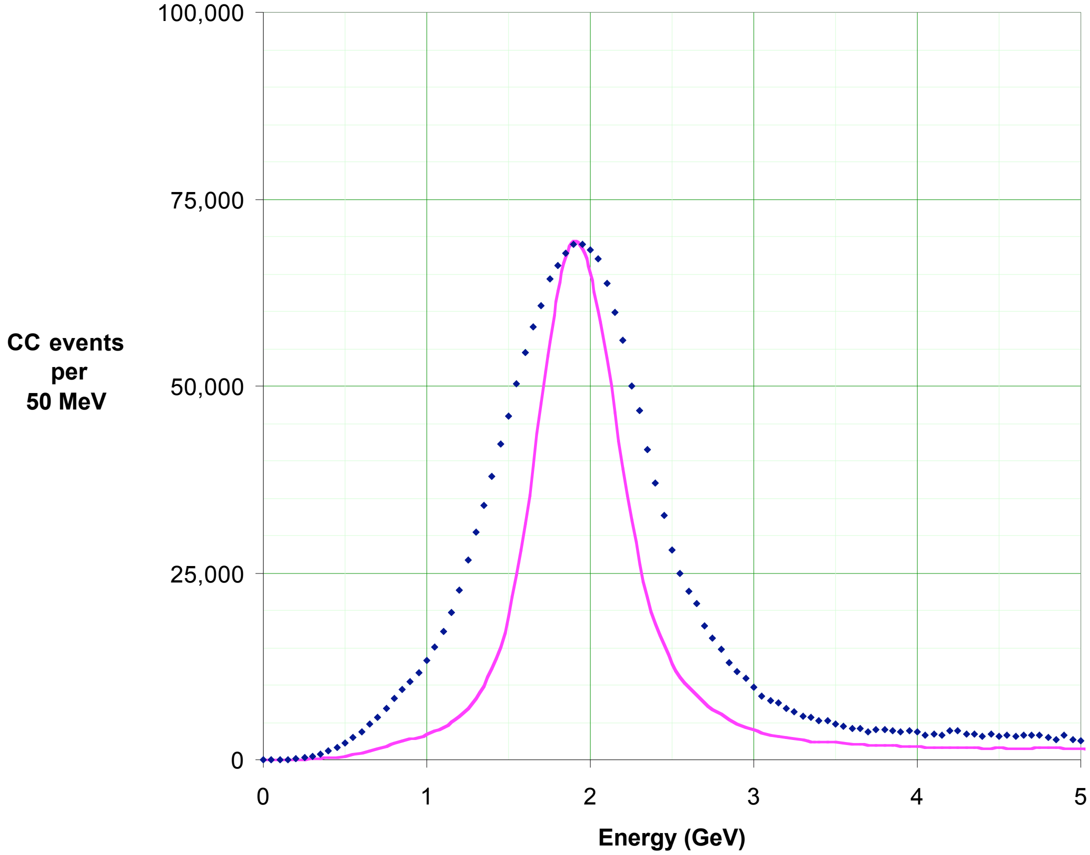

Fig. 1.3: ν_μ CC spectrum for Near Detector Site 1.75 (blue dots) and the unoscillated Far Detector (pink line) for one year of running at 6.5 10^{20} pot. The Far Detector spectrum has been normalized to the Near Detector spectrum at their common peak location.

Question 2: What is your sensitivity (to ν_e appearance and to $\sin^2(2\theta_{13})$ *vs*. calendar time for 20%, 10%, and 5% systematic errors?

Answer to Question 2: The 3-σ sensitivities to ν_e appearance and to $\theta_{13} \neq 0$ are shown in Figures 2.1 and 2.2, respectively, for the parameters specified in the figures. The typical δ is the value of δ that has the median sensitivity. The figures assume that the Project starts in FY2007 and is funded at a rate that allows us to maintain the schedule specified in our proposal.

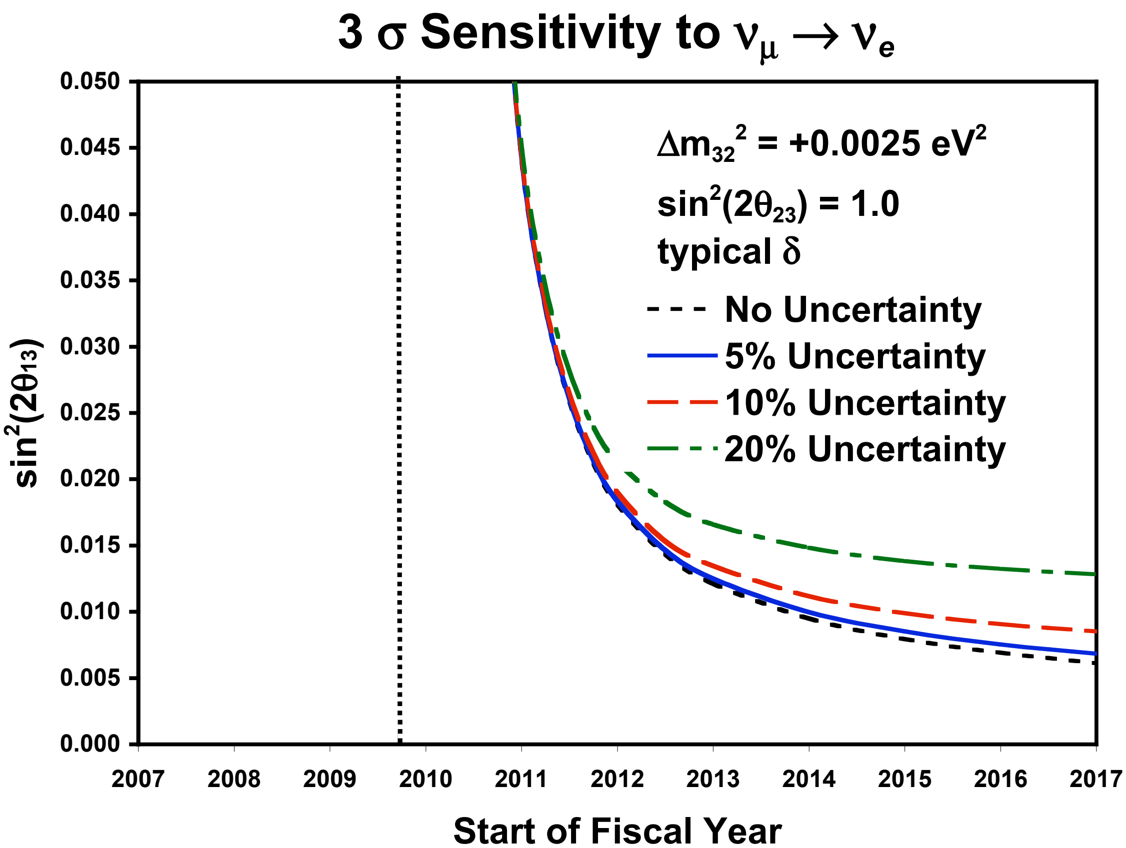

Fig. 2.1: 3-σ sensitivity to ν_e appearance for parameters and systematic uncertainties on the background as specified in the figure. Neutrino only running is assumed. With the exceptions of this figure and the following one, we assume a 5% systematic uncertainty.

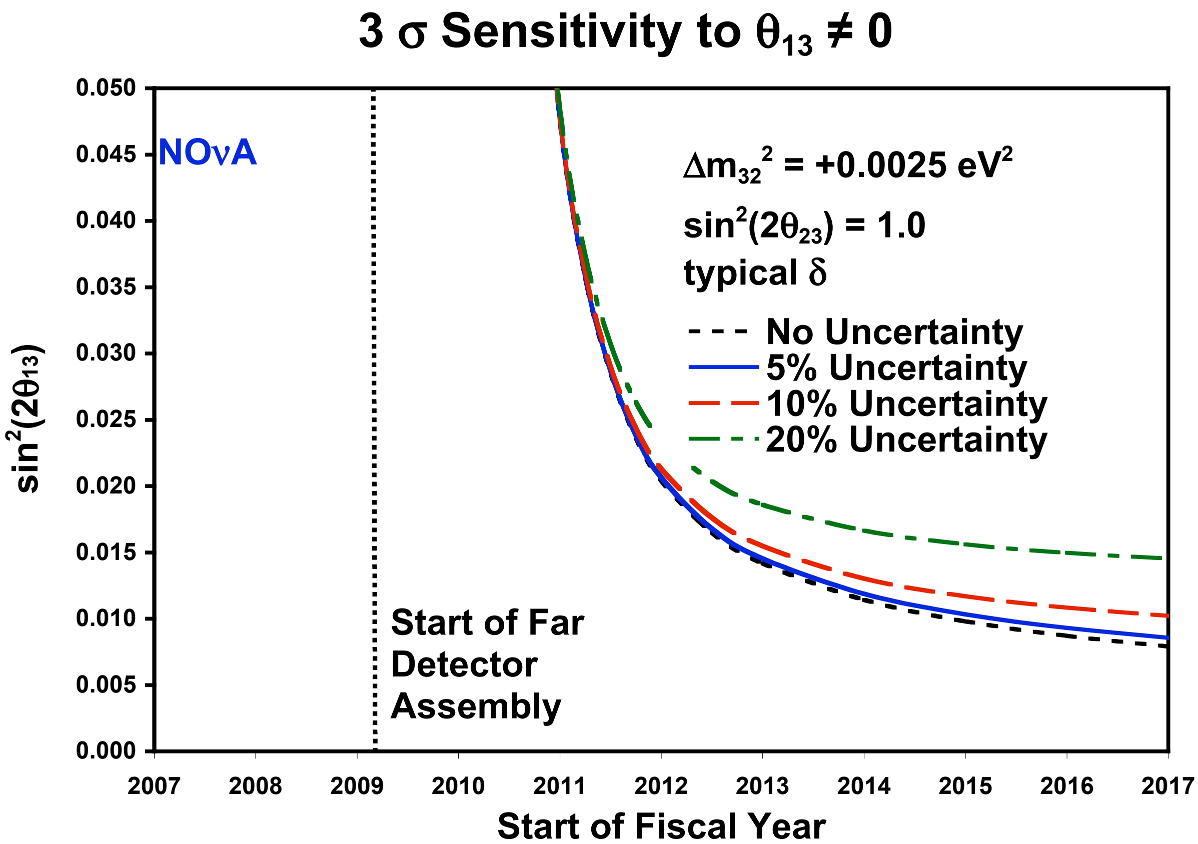

Fig. 2.2: 3-σ sensitivity to $\theta_{13} \neq 0$ for parameters and systematic errors on the background as specified in the figure. Neutrino only running is assumed. With the exception of this figure and the preceding one, we assume a 5% systematic uncertainty.

Question 3: In the absence of a reactor θ_{13} measurement, how well and unambiguously can you determine θ_{13} and θ_{23}?

Answer to Question 3: Accelerator experiments are sensitive, for the most part, to $\sin^2(\theta_{23})\,\sin^2(2\theta_{13})$ and $\sin^2(2\theta_{23})$. The former is measured from $\nu_\mu \to \nu_e$ appearance; the latter is measured from ν_μ disappearance.

We consider θ_{23} first. Figures 3.1 and 3.2 show how well we can measure $\sin^2(2\theta_{23})$ without and with the Proton Driver. Without the Proton Driver, we can measure $\sin^2(2\theta_{23})$ to 0.004 if it is maximal and to 0.01, if it is not, say $\sin^2(2\theta_{23}) = 0.95$. With the Proton Driver, the corresponding 1-σ uncertainties are about 0.002 and 0.005. These values correspond to errors on θ_{23} of 1.8° and 0.7° without the Proton Driver, and 1.3° and 0.35° with the Proton Driver. However, in the non-maximal case there is a two-fold ambiguity corresponding to whether θ_{23} is greater than or less than 45°. The only practical way to resolve this ambiguity is to compare the measurement of $\sin^2(\theta_{23})\,\sin^2(2\theta_{13})$ in accelerator experiments with the measurement of $\sin^2(2\theta_{13})$ in reactor experiments. Figure 3.3 shows the range in $\sin^2(2\theta_{13})$ - $\sin^2(2\theta_{23})$ space for which this is possible. The range is a function of δ, the mass ordering, and the sign of the ambiguity itself, but it is not a strong function of these parameters, so Fig. 3.3 averages over them. The reactor experiment is assumed to have a 1-σ resolution of 0.43%, as specified in the Braidwood presentation.

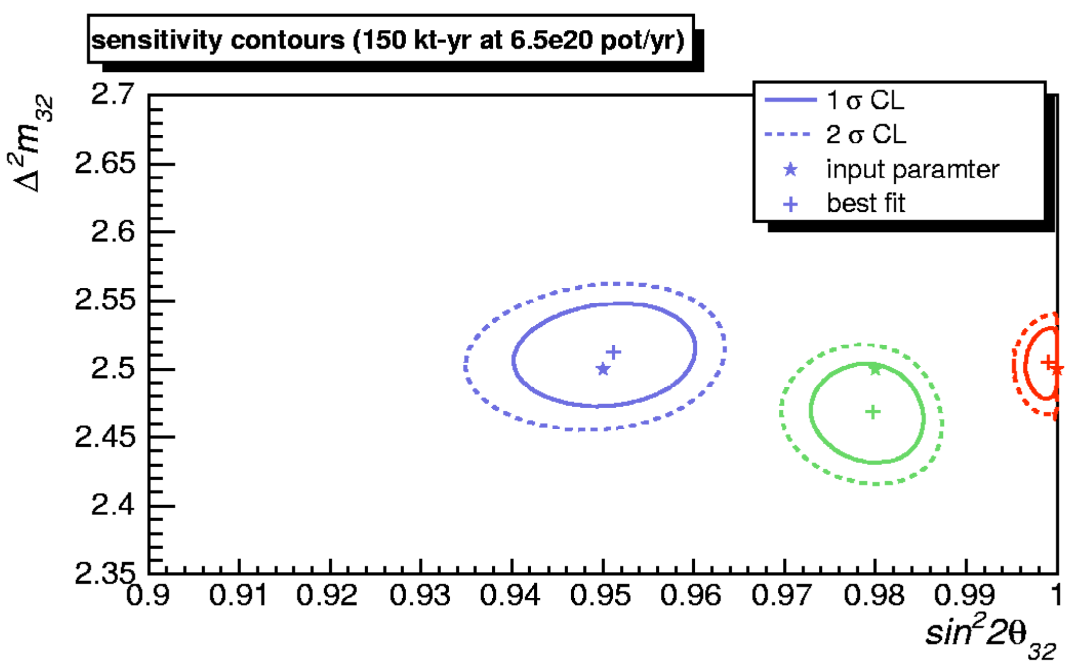

Fig. 3.1: One and two standard deviation contours for the simultaneous measurements of Δm^2_{32} and $\sin^2(2\theta_{23})$ for a five-year neutrino run without a Proton Driver.

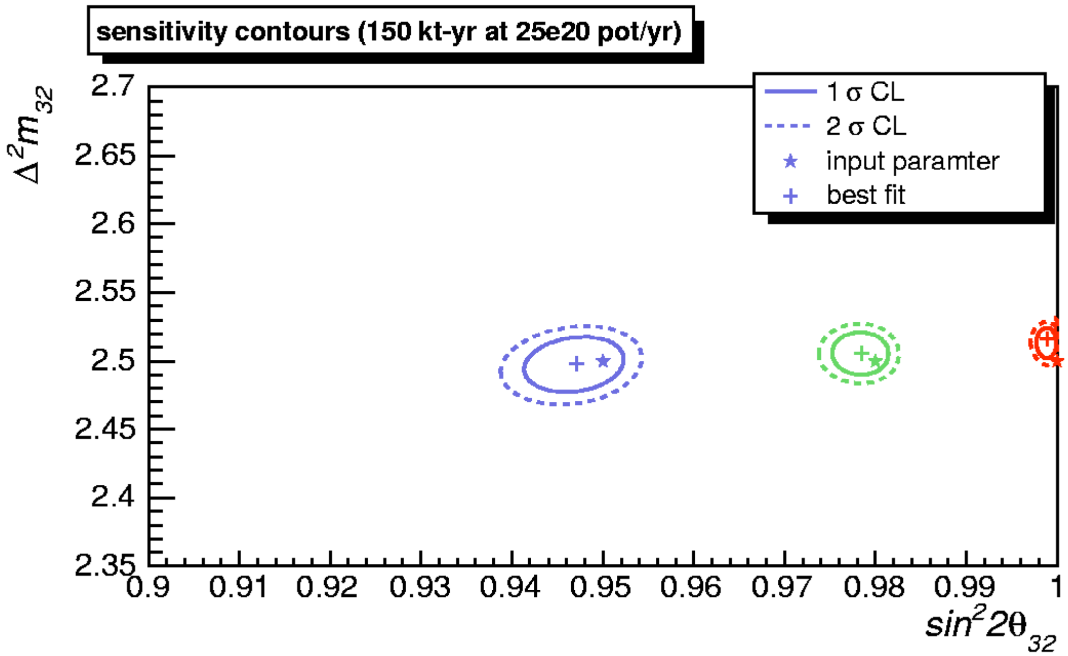

Fig. 3.2: One and two standard deviation contours for the simultaneous measurements of Δm^2_{32} and $\sin^2(2\theta_{23})$ for a five-year neutrino run with a Proton Driver.

We typically emphasize our ability to demonstrate that θ_{13} is non-zero. To do this we need only count Far Detector ν_e CC events in an appropriate energy range, and demonstrate that there is a statistically significant excess of them over what we expect from beam backgrounds and solar-oscillation-length production. However, to measure θ_{13}, we need to perform a simultaneous fit to θ_{13} and δ, for both mass orderings. Examples of the expected results from such fits are shown in Figs. 3.4 through 3.7 for a number of sample points. In these figures, the atmospheric mixing is assumed to be maximal, *i.e.*, $\theta_{23} = 45^{\circ}$. If this is not the case, then, to a reasonably good approximation, the vertical axis scale needs to be divided by $2\sin^2(\theta_{23})$ with the attendant errors and ambiguity discussed above.

Figure 3.4 shows the results for $\sin^2(2\theta_{13}) = 0.05$, $\delta = 270^{\circ}$, normal mass ordering. This is a point at which the mass ordering has been resolved. For contrast, Fig. 3.5 shows the results for $\sin^2(2\theta_{13}) = 0.05$, $\delta = 180^{\circ}$, inverted mass ordering, a point for which the mass ordering is not resolved. Figure 3.6 shows the results for the same parameters as Fig. 3.4, except that $\sin^2(2\theta_{13})$ is reduced from 0.05 to 0.02. Finally, Fig. 3.7 shows the results for the same parameters as in Fig. 3.6, but with a six-year run with the Proton Driver.

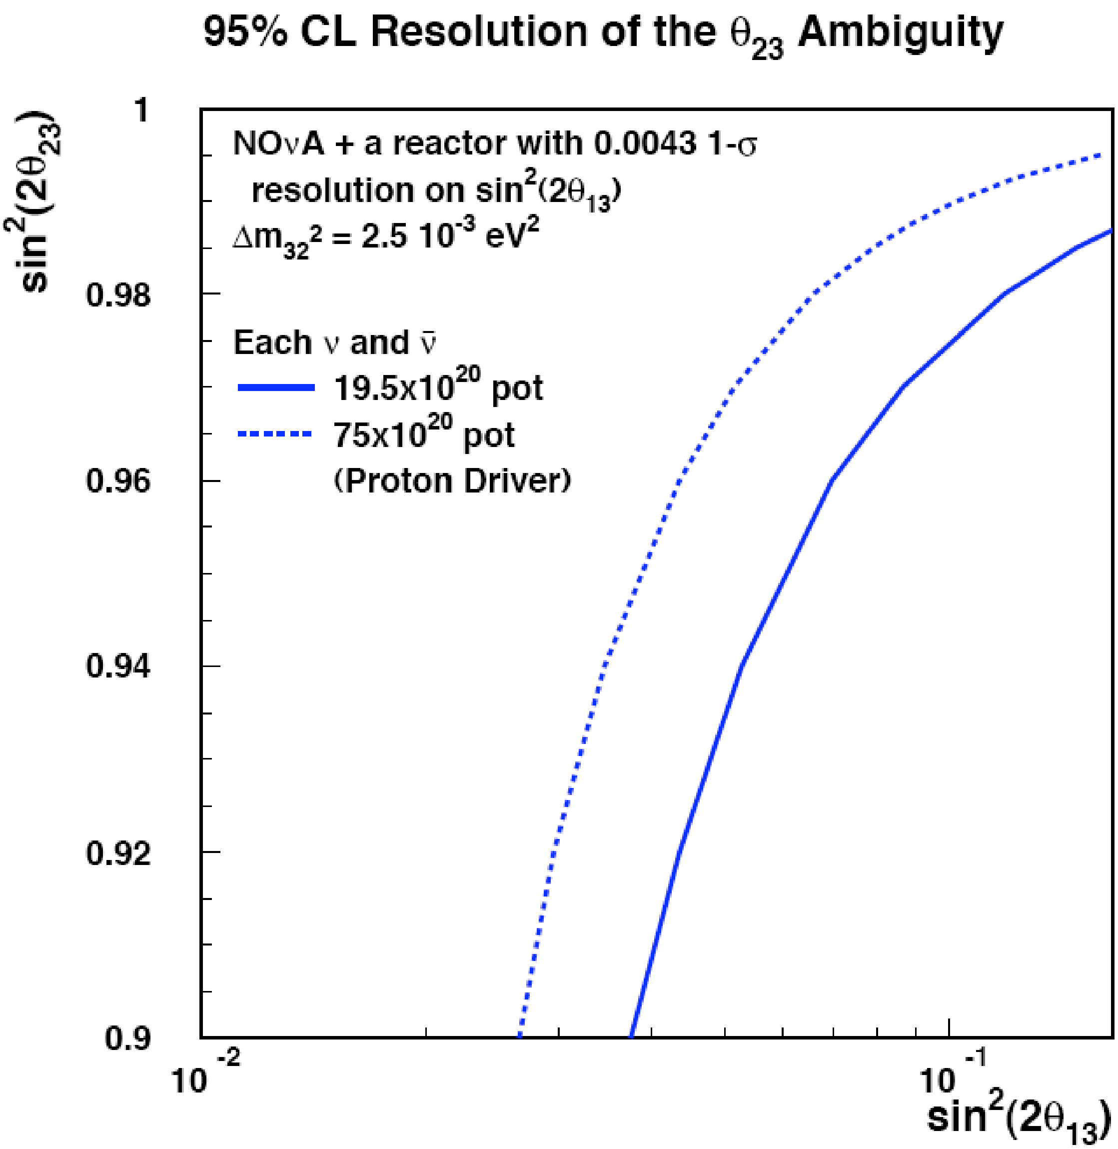

Fig. 3.3: 95% confidence level resolution of the θ_{23} ambiguity. The regions to the right of the curves represent the parameter space for which the ambiguity can be resolved by a six-year NOνA run split evenly between neutrino and antineutrino running and a reactor experiment with 0.43% 1-σ resolution on $\sin^2(2\theta_{13})$. The dotted and solid curves represent the NOνA sensitivity with and without the Proton Driver, respectively. The curves are averaged over the other parameters and they assume that the value of $\sin^2(2\theta_{23})$ is known to high precision.

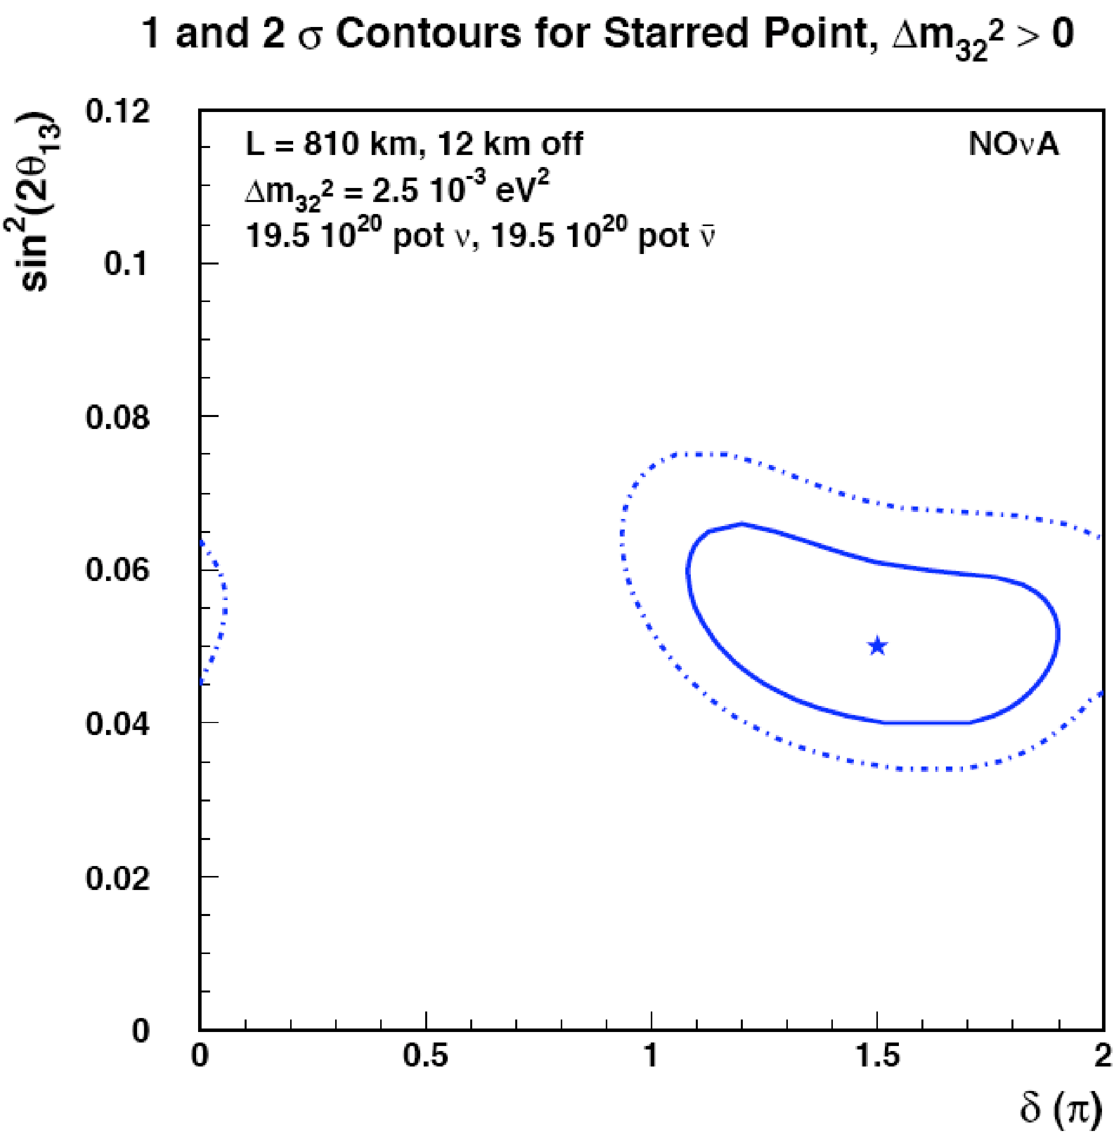

Fig. 3.4: One and two σ contours for a simultaneous fit to $\sin^2(2\theta_{13})$ and δ for the starred point, $\sin^2(2\theta_{13}) = 0.05$, $\delta = 270^o$, normal mass ordering, under the assumption that $\sin^2(2\theta_{23}) = 1$. Only contours for the normal mass ordering are present since the mass ordering is resolved at the 95% confidence level for this point. The figure assumes a six-year NOνA run split evenly between neutrino and antineutrino running.

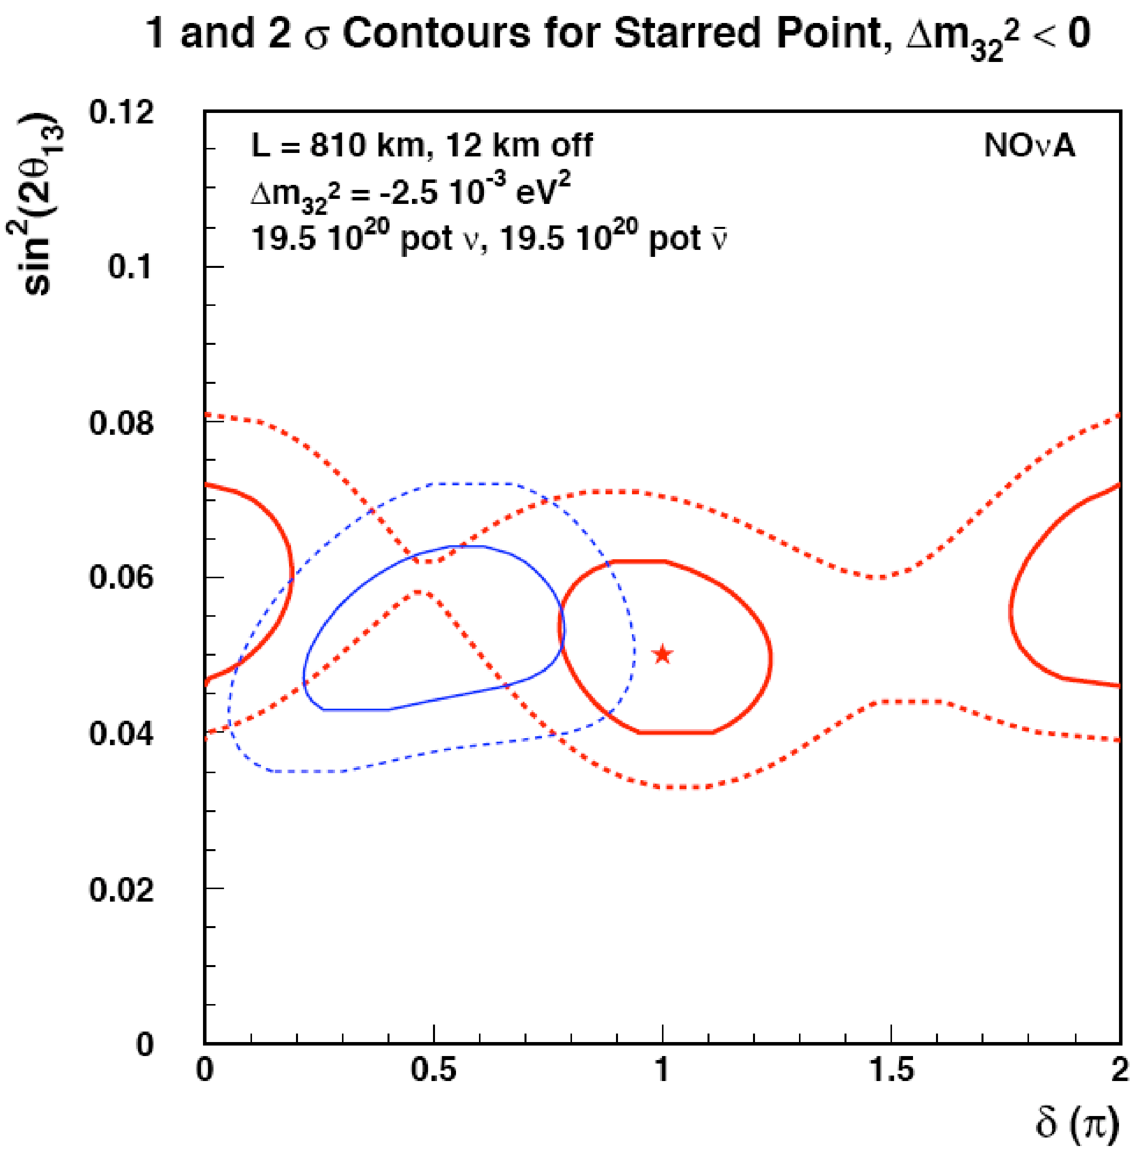

Fig. 3.5: One and two σ contours for a simultaneous fit to $\sin^2(2\theta_{13})$ and δ for the starred point, $\sin^2(2\theta_{13}) = 0.05$, $\delta = 180^o$, inverted mass ordering, under the assumption that $\sin^2(2\theta_{23}) = 1$. The red contours are for the inverted mass ordering and the blue contours are for the normal mass ordering. The figure assumes a six-year NOνA run split evenly between neutrino and antineutrino running.

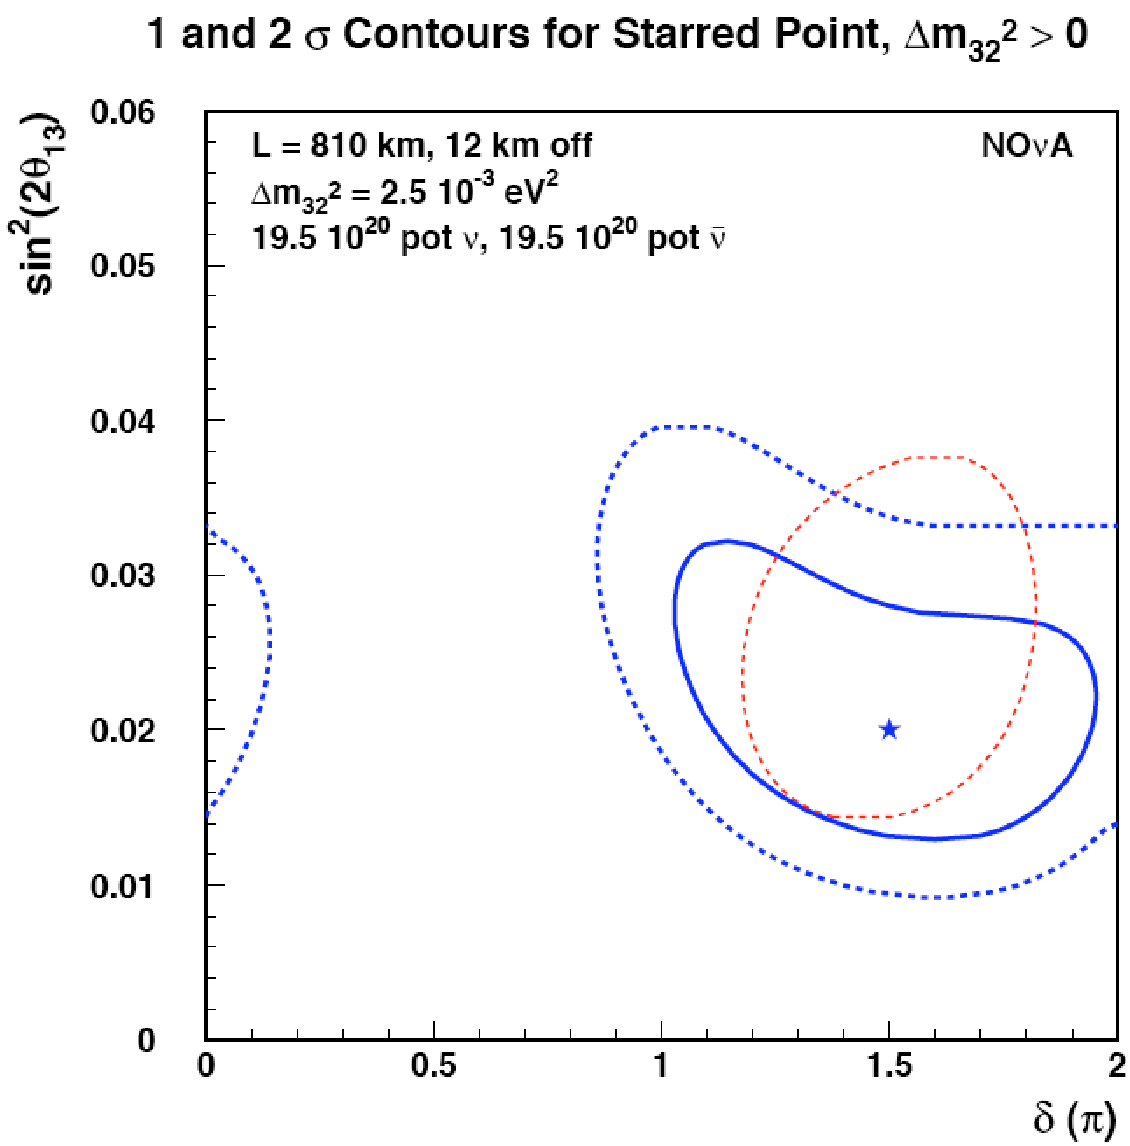

Fig. 3.6: One and two σ contours for a simultaneous fit to $\sin^2(2\theta_{13})$ and δ for the starred point, $\sin^2(2\theta_{13}) = 0.02$, $\delta = 270^o$, normal mass ordering, under the assumption that $\sin^2(2\theta_{23}) = 1$. The blue contours are for the normal mass ordering and the red contour is for the inverted mass ordering. The 1-σ inverted mass ordering contour is not present because the mass ordering is resolved to that level for this point. The figure assumes a six-year NOνA run split evenly between neutrino and antineutrino running.

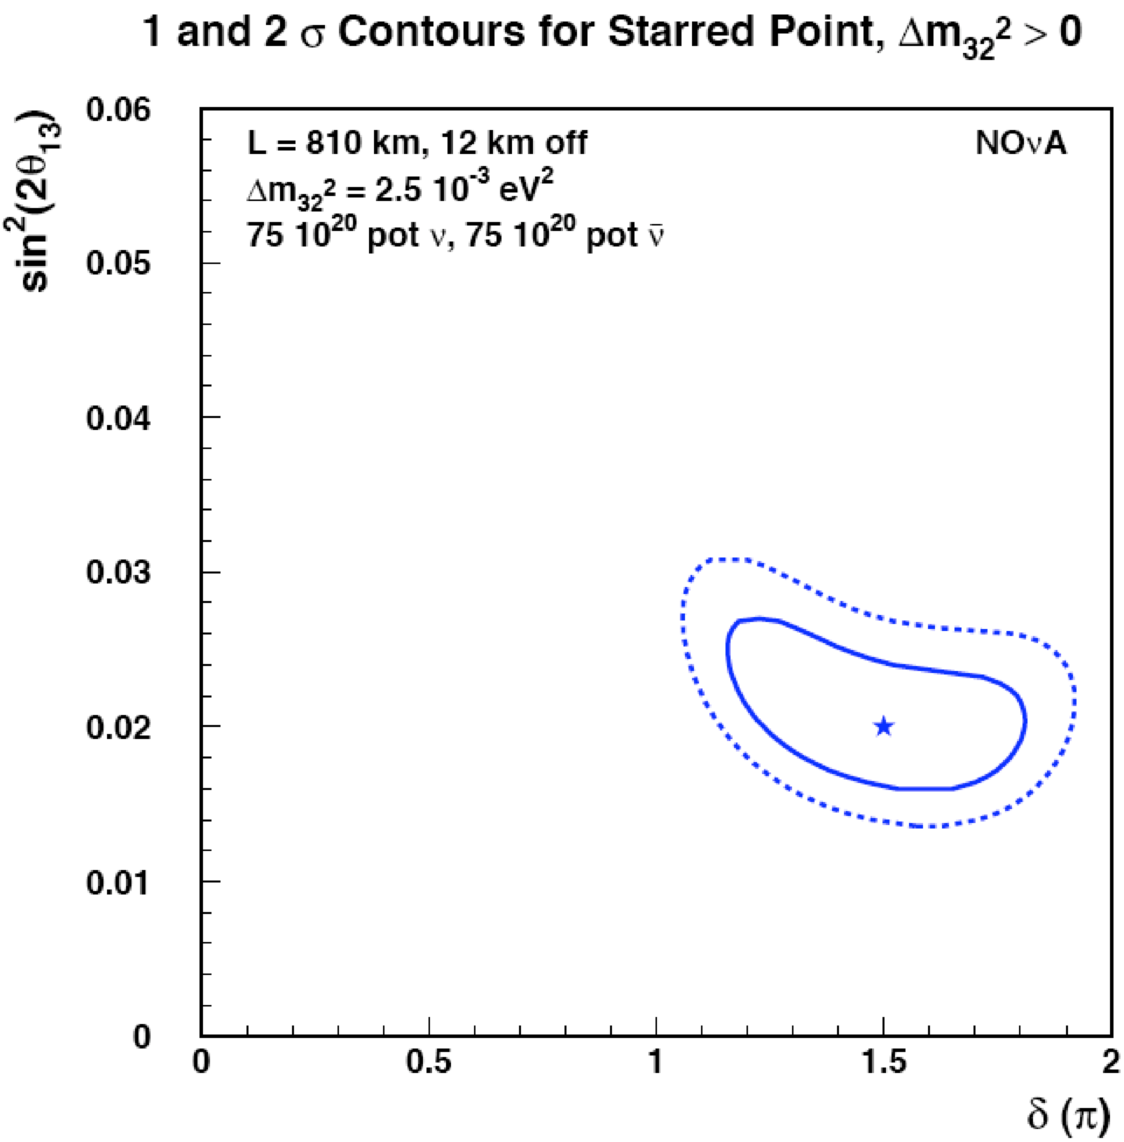

Fig. 3.7: One and two σ contours for a simultaneous fit to $\sin^2(2\theta_{13})$ and δ for the same parameters as in Fig. 3.6, $\sin^2(2\theta_{13}) = 0.02$, $\delta = 270^o$, normal mass ordering, with a six-year NOνA run with the Proton Driver. Only contours for the normal mass ordering are present because the mass ordering has been resolved to the 95% confidence level for this point.

Question 4: What is the effect of Δm^2_{32} on your sensitivity?

Answer to Question 4: Figure 4.1 shows the 3-σ sensitivity to $\theta_{13} \neq 0$ as a function of Δm^2_{32} for a six-year NOνA run split evenly between neutrino and anti-neutrino running for the typical value of δ, i.e., the value of δ that gives the median sensitivity.

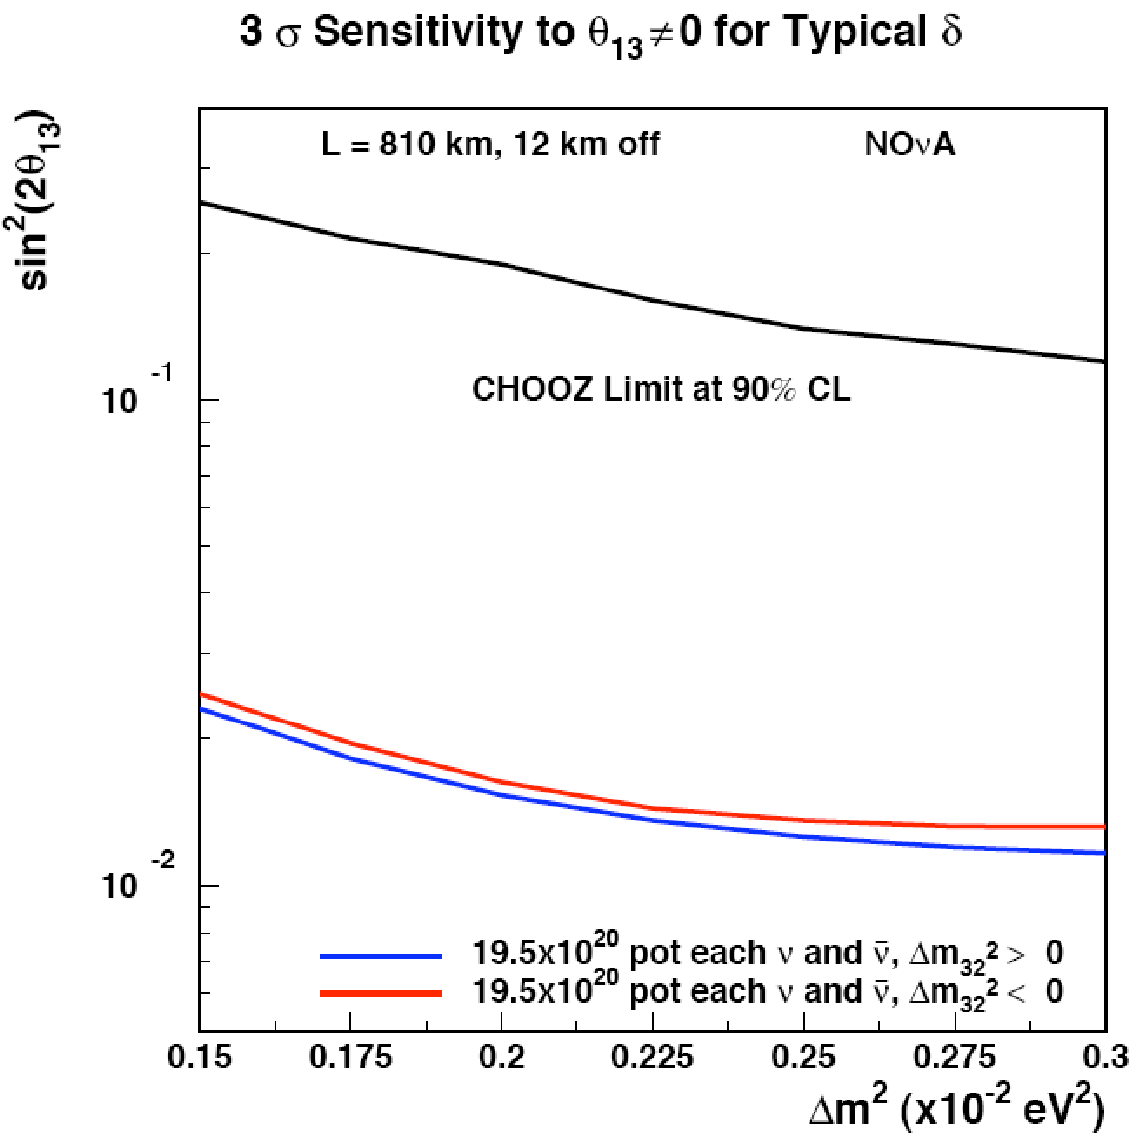

Fig. 4.1: 3-σ sensitivity to $\theta_{13} \neq 0$ as a function of Δm^2_{32} for a six-year NOνA run split evenly between neutrino and anti-neutrino running for the typical value of δ. The blue and red curves represent the normal and inverted mass orderings, respectively. The 90% confidence level CHOOZ limit is shown for comparison.

Question 5: How much would improved neutrino cross section measurements, using appropriate targets for your experiment, allow you to reduce the systematic uncertainty on your background estimate? In particular, please evaluate the potential that the MINERνA experiment has to improve your result, and whether or not you are counting on having these measurements.

Answer to Question 5: MINERνA will be a fine-grained detector experiment, which should be able to measure total cross sections, different channel cross sections, and angular distributions with a precision limited only by knowledge of the flux. These detailed measurements will undoubtedly be valuable in tuning our Monte Carlo simulations, and, as stated in our answer to Question 1, we need to perform a careful study to be able to quantify this.

However, we note that the information we need to measure backgrounds and interpret signals is the product of our experimental efficiency times the cross section as a function of visible energy for the three general classes of events, ν_e CC events, ν_μ CC events, and NC events. The best and most direct source of this information will be the MIPP experiment's determination of neutrino fluxes combined with our measurements in the Near Detector. Breakdowns of the cross sections into the traditional categories of quasielastic, resonance, and deeply inelastic scattering are not needed for this purpose.

Question 6: Figure 13.27 in the proposal shows the calculated ν_μ contamination in your $\bar{\nu}_\mu$ beam. How well does it have to be known to reach the combined $\nu + \bar{\nu}$ results you quote? How will this be measured? How well must the $\bar{\nu}_\mu$ cross sections be known? Some of the backgrounds come from feed-down from higher energies, where the ν_μ contamination is ~100%. Is this included in the sensitivity estimates in Ch. 13?

Answer to Question 6: First the numbers: At the Far Detector, integrated over our energy acceptance, and assuming equal oscillation probabilities for neutrinos and antineutrinos, in neutrino running, 0.9% of the observed signal events will be from antineutrinos, and in antineutrino running, 6.7% of the observed signal events will be from neutrinos. This includes feed-down from higher energies, which is negligible due to both the relatively low number of higher energy events and the excellent energy resolution of the detector.

The effect of "wrong sign" neutrinos depends on what parameter is being measured. For detection of the signal, *i.e.*, sensitivity to $\theta_{13} \neq 0$, the results of neutrino and antineutrino running are just added together. Thus, the small fraction of wrong sign neutrinos is equivalent to a slight modification of the relative amount of neutrino and antineutrino running, and does not contribute any significant additional error. However, for determination of the mass ordering and CP violation, which requires a measurement of the ratio of neutrino to antineutrino oscillation probabilities, the wrong sign events degrade the statistical precision by between 4 and 14%, depending on the values of the physics parameters. This was not included in the sensitivity estimates in Chapter 13.

Both antineutrino fluxes and cross sections will be needed to calculate the neutrino background in the antineutrino beam and, more importantly, to interpret signals in antineutrino running. The fluxes will be measured in the MIPP experiment. The best way to measure the antineutrino cross sections is with a magnetic analysis. This will be done eventually in the MINOS Near Detector and in the MINERνA experiment using the MINOS Near Detector as a muon spectrometer. However, a magnetic analysis is not necessary as long as the fluxes are known, so the NOνA Near Detector can also make an important contribution. Given the expected statistical errors, a 10% measurement of the antineutrino cross sections in the 2 GeV energy region should be sufficient, and this should be easily obtainable in these experiments.

Question 7: In the proposal you state that you will have new production measurements from E907 the MIPP experiment. Specifically what measurements are needed?

Answer to Question 7: The MIPP measurements are extremely important for our measurements since they will provide the data from which we can calculate the ν_μ, $\bar{\nu}_\mu$, ν_e, and $\bar{\nu}_e$ fluxes at both the Near and Far Detectors. These measurements will be particularly useful since MIPP is an open geometry experiment. MIPP is taking data on the actual (spare) NuMI target. All of the particles that we are interested in for neutrino production will be seen and momentum analyzed in every event and they will all be identified as pions, kaons, or protons by the MIPP Cerenkov counters. Thus, MIPP removes the need for a Monte Carlo simulation of hadron production on the NuMI target, since a library of corrected MIPP events can be used instead. MIPP hopes to measure hadron production cross sections to the 5% level, many of which are now only known to 20%.

Another important contribution of the MIPP flux determination will be in determining the energy scale, i.e., the conversion of visible energy to true energy. Invariably, when a neutrino undergoes an interaction, some energy is lost or mismeasured due to incalculable process such as nuclear scattering, very soft fragmentation, and low-energy neutron emission. Given sufficiently accurate MIPP data, we can very precisely calculate the peak energy for ν_μ CC events in the Near Detector (see Fig. 1.3) and use it to calibrate our visible energy measurement.

Unfortunately, the number of MIPP events with the NuMI target will be substantially smaller than we would like. The original goal was 10 million events on the NuMI target. As of this writing, it appears that this year's MIPP run will produce only 10 to 20% of that number. We support additional MIPP running.

Question 8: The light collection in the near and far detectors is different, how will this affect the relative energy scale and resolution in the detectors?

Answer to Question 8: The energy scale will not be affected since it is calibrated by using stopping cosmic muons, as in MINOS.

Our light requirement is set by pattern recognition, not by resolution. We want to be able to measure a fraction of the output of a minimum-ionizing particle (mip), and, for this reason, we require a minimum signal-to-noise ratio of 10 from the far end of each cell. We plan to alternate the location of readouts for the horizontal modules. For a typical 2 GeV event, this design leads to a signal of between roughly 6000 and 18000 photoelectrons, depending on the location of the event in the Far Detector, over a maximal electronic noise of 35 photoelectrons. Thus, it is clear that our resolution of approximately $10\% / \sqrt{E}$ is dominated by event fluctuations rather than photon statistics.

In the answer to Question 11, we will discuss decreasing the shaping time and increasing the sampling frequency in the Near Detector to gain better timing resolution. The only effect on energy resolution of this difference will be to increase the noise level in the Near Detector. From the numbers presented above, this is clearly not a concern.

We also note that, if necessary, it is relatively trivial in software to degrade Near Detector events so that they simulate events in different parts of the Far Detector.

Question 9: In order to reduce the ν_e contamination in the Far Detector you make an energy cut on the reconstructed energy. Can you be more specific about the cut, the reduction it gives, and the effect on the systematic error for energy scale and resolution differences in the Near and Far detectors?

Answer to Question 9: One of the primary advantages of an off-axis beam is that the signal is concentrated in a narrow energy band while the background is likely to be outside that band. This can be seen in Figure 9.1. The NC events are strongly peaked at lower visible energies than the signal, since the neutrino carries away some of the energy. The beam ν_e events have sizable contributions both above and below the signal region because they come from muon decays, which are three-body, and from kaon decays, which can emit leptons at higher energy into the angular acceptance of the detector.

In the analysis presented in the proposal, a preliminary cut on visible energy was made between 1.5 and 2.5 GeV since the signal drops approximately an order of magnitude beyond these limits. Then, the visible energy was used as an explicit parameter in the construction of the likelihood function. The preliminary cut was made mostly for convenience. Presumably the likelihood function would have rejected events beyond these limits, so the analysis is likely to have been almost the same if the preliminary cut had not been made.

There should be no large systematic uncertainty associated with the use of this variable since the Near and Far detectors will be cross-calibrated with stopping cosmic rays, as explained in the answer to Question 8. The philosophy of the analysis is to treat the Near and Far detectors identically. Of course, there will be some unavoidable differences between the detectors, and the consequences of these differences need to be simulated and corrected for in the analysis, and an appropriate systematic error assigned. Issues of containment, data rate, and non-beam backgrounds are likely to be more of a concern than the energy variable.

From Fig. 9.1, one can assume that the reduction in background due to the use of the visible energy variable is quite large. How large is not well defined since no sensible analysis would admit events from regions where the likelihood of finding signal events is very small.

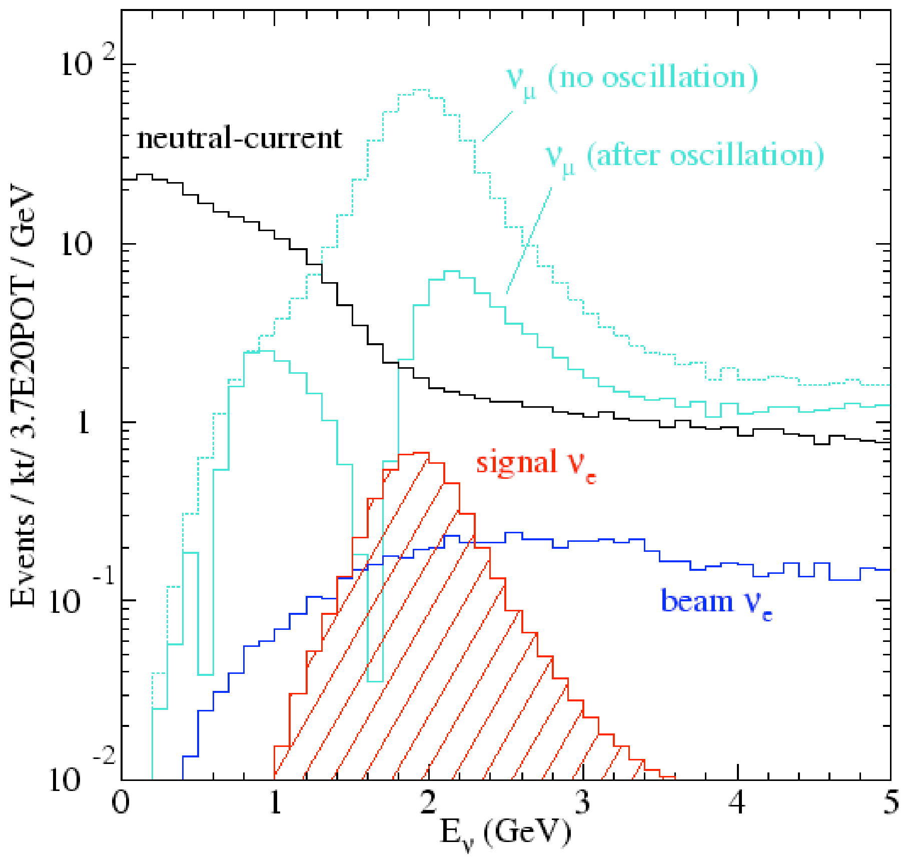

Fig. 9.1: Simulated energy distributions for the ν_e oscillation signal, intrinsic beam ν_e events, neutral-current events and ν_μ charged-current events with and without oscillations.

Question 10: In order to accurately measure your ν_e and π^0 backgrounds and to minimize systematic error, your plan calls for moving the near detector several times (*e.g.,* test beam, multiple near locations, the MINOS surface building). Has a detailed execution plan and approximate cost been developed for this? Is it included in the cost estimate and schedule?

Answer to Question 10: Our present plans are to have both a Prototype Near Detector and a Near Detector. The Prototype Near Detector will be built on R&D funds and is scheduled to be completed by March 2007. Although we have not worked out a detailed execution plan, it would probably be sensible to have the Prototype Near Detector in the test beam and the MINOS Surface Building and the Near Detector underground. (This is similar to what MINOS did; the Calibration Detector was put in a test beam at CERN.) Based on preliminary design considerations, we have budgeted $578k for the carriage to move the Near Detector in the underground NuMI tunnel. This cost includes a large contingency. We anticipate that the Near Detector would be moved infrequently, perhaps once or twice a year.

Question 11: How will the increased number of multiple overlapping events in the near detector with a Proton Driver affect the background measurements?

Answer to Question 11: Overlapping events in the Near Detector are a concern with or without the Proton Driver. Before the Proton Driver we expect 3.8 events per 10 μs spill in the Near Detector, and this rate will increase to 9.5 events per spill with the Proton Driver. (The remainder of the rate increase with the Proton Driver comes from decreasing the Main Injector cycle time.) We will use the same technique to separate events that MINOS does, by using position and timing information.

We need to perform some simulations to determine the overlap probability of events in space. We note that the worst case for a 2-GeV event is that it will occupy about 120 cells, or 1% of the cells in the Near Detector.

There are two issues with regard to time separation, the time resolution of a cell that is hit by only one event, and the minimum time separation to separate two events that hit the same cell. In the Far Detector we plan to use electronics that has an approximate shaping time of 200 ns and a sampling frequency of 2 MHz. We have simulated the timing resolution for the case in which only one event hits a cell. The result is shown in Fig. 11.l. Depending on the location in the Far Detector, the rms time resolution will be between 20 and 60 ns. A conservative estimate of the two-event separation resolution is five sampling times, or 2.5 μs.

Since there will be a minimum of 50 photoelectrons from a 1-mip signal in the Near Detector, we can afford to reduce the shaping time and increase the sampling frequency in the Near Detector by up to a factor of 4 and still maintain our required minimum signal-to-noise ratio of 10. Scaling the result from Fig. 11.1, the maximum rms timing resolution will be 15 ns and the conservative estimate of the two-event separation resolution will be 625 ns.

We are also considering instrumenting several planes in the Near Detector with MINOS detectors (multi-anode photomultipliers) and electronics. The MINOS Near Detector electronics samples pulse heights every 19 ns. This will allow us to reduce the two-event separation resolution to about 100 ns, and, at a minimum, alert us to an overlap situation.

In cases in which the temporal and spatial overlap between events makes it impossible to disentangle them, the overlapping events will be discarded. There is some bias in this procedure, since the special overlap probability will be a function of the event size. This bias needs to be simulated and corrected for. These corrections can be checked by special low intensity runs in which the expected number of events is one per spill.

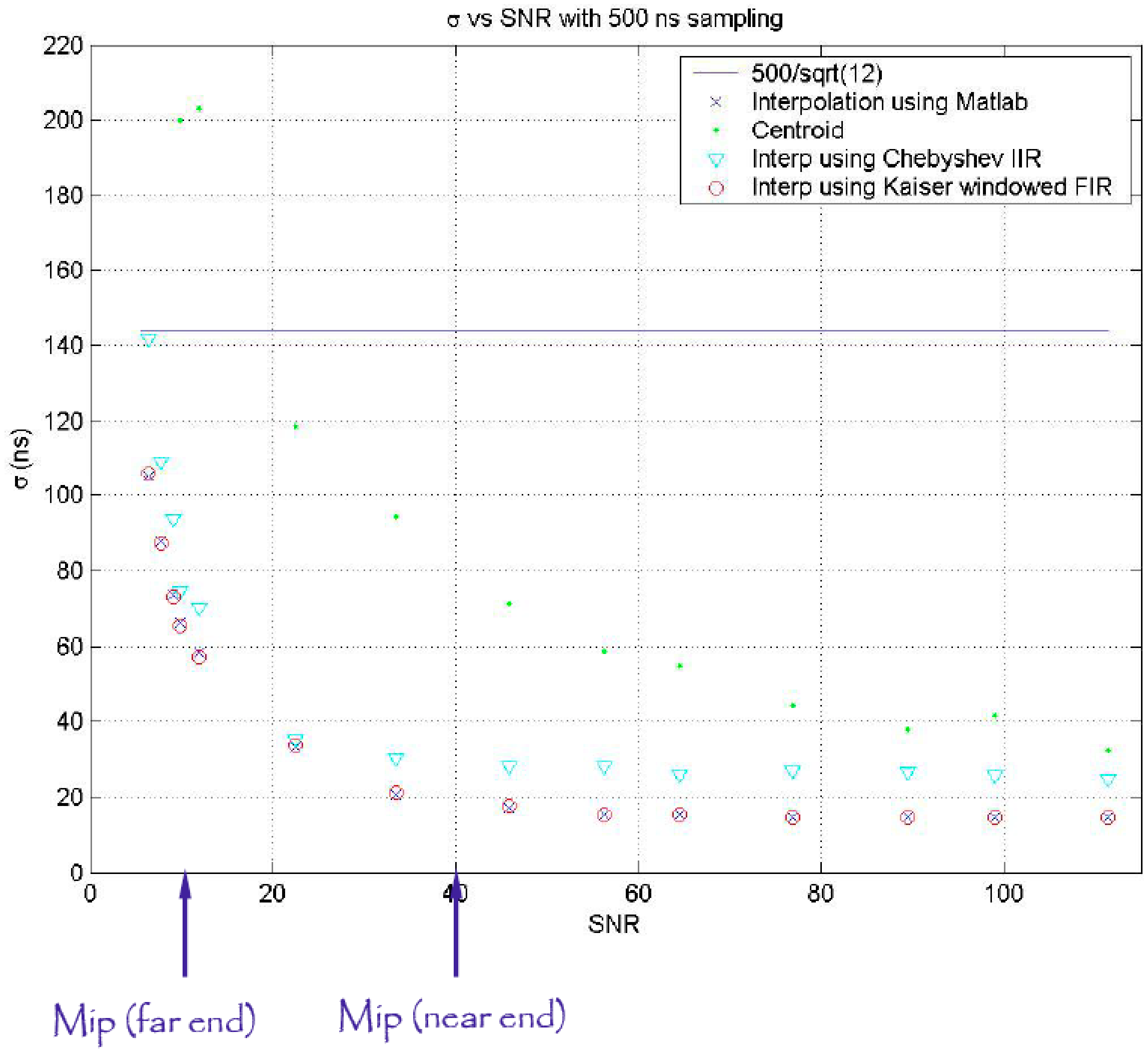

Fig. 11.1: Rms time resolution for a cell hit by only one event as a function of signal-to-noise ratio (SNR), using several different algorithms. The shaping time is assumed to be approximately 200 ns and the sampling frequency is assumed to be 2 MHz. The best results are obtained with a Kaiser windowed finite impulse response filter, shown in the open circles.

Question 12: On page 71 of the proposal you list the estimated backgrounds for 5 years of running to be 11.9 events from intrinsic ν_e, 0.5 events from ν_μ CC, and 7.1 events from NC π^0 production. These will be known to levels estimated to be 7%, 15%, and 5%, respectively. What measurements do you make in order to reduce the π^0 background uncertainty to 5%?

Answer to Question 12: The data from the Near Detector in the MINOS tunnel and from Prototype Near Detector in MINOS Surface Building will be very useful in refining estimates of π^0 background. Most of the photons from π^0 s in NOνA detectors will convert sufficiently far from the neutrino interaction vertex so that two photons can be clearly identified. We can study what fraction of these events would simulate ν_e CC events if conversion would occur near the vertex by moving the conversion point in software.

See the answers to Questions 1, 5, and 7 for more details on our approach to measuring backgrounds.

Question 13: Your calculations indicate that cosmic-induced backgrounds will not be a problem for the ν_e search. Is this also the case for the other physics goals, *e.g.,* ν_μ disappearance? On page 53 of the proposal, the possibility of an overburden to mitigate cosmics is discussed. What is a rough estimate of the additional cost of this option?

Answer to Question 13: Section 10.9 of our proposal does state that cosmic ray backgrounds "should not be a problem." Unfortunately, after the proposal was written, we discovered that this is not completely correct. We have convinced ourselves that all sources of cosmic ray backgrounds are not a problem except for the photon component, which can simulate ν_e CC events at an unacceptable rate. We can reduce the level of this problem by tuning our analysis to better reject this background, and we are presently studying this possibility. However, it is possible that this background will remain unacceptable for a completely unshielded Far Detector and for this reason we are considering the option of a three-meter rock overburden for the detector, using rock excavated during construction of the Detector Hall. This option has a number of other advantages, including a much improved supernova signal and reduced requirements for DAQ and HVAC systems. Our present estimate for the additional cost of a three-meter overburden is $ 7.6M, without contingency.

There are no sources of cosmic background that can simulate contained ν_μ CC events in the Far Detector at a significant rate.